

Trustworthiness as a Limitation on Network Neutrality*

Aaron J. Burstein[†] and Fred B. Schneider[‡]

Introduction	3
I. Tracing Trustworthiness Through the Network Neutrality Debate.....	5
A. Defining Trustworthiness	5
B. Trustworthiness as a Limitation on Nondiscrimination in Common Carrier Regulations	10
C. Trustworthiness as a Limitation in Network Neutrality Rules.....	15
II. The Implications of Trustworthiness Exceptions in Network Neutrality Rules	21
A. Isolation from Unwanted Traffic.....	22
1. Isolation as a Consumer Service	23
2. Isolation as Provider Policy	24
3. Isolation Under the Broad Exception	27
B. Availability and Integrity: Attribution of Path.....	28
C. Privacy and Confidentiality: Guarantees Against Logging	34
D. Trustworthiness and Wireless Net Neutrality	38
III. Keeping Trustworthiness Exceptions Limited Through Disclosure.....	40
A. Mechanisms to Deter Trustworthiness-as-Pretext.....	40
B. Striking the Right Balance for Trustworthiness Disclosures	43
Conclusion.....	46

* A preliminary version of this article was presented at the 35th Research Conference on Communication, Information and Internet Policy (TPRC), Arlington, VA, September 28, 2007. The authors acknowledge helpful comments from TPRC participants. The authors also acknowledge helpful comments from Deirdre Mulligan on earlier drafts of this article.

[†] TRUST and ACCURATE Research Fellow, Samuelson Law, Technology & Public Policy Clinic and Berkeley Center for Law and Technology, School of Law, University of California, Berkeley. Burstein acknowledges support for this work from TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422).

[‡] Professor, Department of Computer Science, Cornell University. Schneider acknowledges support from TRUST, AFOSR grant F9550-06-0019, NSF grant 0430161, and funding from Microsoft Corporation.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2007		2. REPORT TYPE		3. DATES COVERED 00-00-2007 to 00-00-2007	
4. TITLE AND SUBTITLE Trustworthiness as a Limitation on Network Neutrality				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Cornell University, Department of Computer Science, Ithaca, NY, 14853				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Accepted for publication, Federal Communications Law Journal, Vol. 61. U.S. Government or Federal Rights License					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 47	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Abstract

The policy debate over how to govern access to broadband networks has largely ignored the objective of network trustworthiness—a set of properties (including security, survivability, and safety) that guarantee expected behavior. Instead, the terms of the network access debate have focused on whether imposing a nondiscrimination or “network neutrality” obligation on network providers is justified by the condition of competition among last-mile providers. Rules proposed by scholars and policymakers would allow network providers to deviate from network neutrality to protect network trustworthiness, but none of these proposals has explored the implications of such exceptions for either neutrality or trustworthiness.

This article examines the relationship between network trustworthiness and network neutrality and finds that providing a trustworthiness exception is a viable way to accommodate trustworthiness within a network neutrality rule. Network providers need leeway to block or degrade traffic within their own subnets, and trustworthiness exceptions can provide them with sufficient flexibility to do so. But, the article argues, defining the scope of a trustworthiness exception is critically important to the network neutrality rule as a whole: an unduly narrow exception that could thwart innovative network defenses, while a broad exception could allow trustworthiness to become a pretext that protects a wide range of discrimination that network neutrality advocates seek to prevent. Furthermore, monitoring network providers’ use of a trustworthiness exception is necessary to ensure that it remains an exception, rather than becoming a rule. The article therefore proposes to require network providers disclose data regarding their

use of a trustworthiness exception and offers a general structure for managing these disclosures.

Introduction

In the United States and other technologically advanced countries, individuals, businesses, and governments have come to depend on the Internet. Daily reports of attacks, accidental data leaks, and service disruptions suggest that the proper functioning of the Internet is not something to take for granted. Trustworthiness—a concept that encompasses not only security but also safety, survivability, and other properties that guarantee expected behavior—is becoming a prominent research and public policy objective.

Internet trustworthiness is hardly the only objective of Internet policy, and setting the terms under which new applications and content sources can reach Internet users has become a focus of much recent debate. Scholars and policymakers have cast this debate in terms of the network neutrality, which holds that network providers may not block, degrade, or otherwise discriminate against applications or content sources. A permissive regulatory environment might allow such discrimination, and that the lack of competition in last mile broadband connections might well make it profitable.

What are the implications of a network neutrality rule for trustworthiness (and vice versa)? Scholars and policymakers have thus far given only superficial answers to this question or avoided it entirely, concentrating instead on whether a network neutrality rule would help or hurt innovation on the Internet. Network neutrality opponents argue that improved security is one type of innovation that might follow from not imposing a network neutrality rule, but this ignores the technical and economic issues that make

improving trustworthiness a hard problem.¹ Proponents, on the other hand, concede that network security is crucial enough to warrant making exceptions to a network neutrality rule. Allowing network providers to deviate from neutrality only to the extent necessary to protect network trustworthiness is rooted in judicial and regulatory decisions and administrative rules that helped establish nondiscrimination as the core of network neutrality. This doctrine of trustworthiness-by-exception stretches back over 50 years and developed around the telephone network. Whether this doctrine is suitable for the technical and institutional complexity of the Internet is unclear, and network neutrality proponents have not made the case that it applies.

We argue in this article that using trustworthiness as a limitation on network providers' nondiscrimination obligations is basically sound and that the set of trustworthiness mechanisms network operators may deploy depends heavily on the exact language of the (proposed) exception.² Some existing proposals would likely thwart valuable trustworthiness mechanisms while others could allow network providers to use trustworthiness as a pretext to discriminate while doing little to improve trustworthiness. Still, there is a middle ground that accommodates neutrality and trustworthiness.

The article is structured in three parts. Part I defines trustworthiness and shows that it has served as a limitation on network operators' nondiscrimination obligations throughout the development of network competition policy and scholarship. Reviewing current

¹ See, e.g., Christopher S. Yoo, *Beyond Network Neutrality*, 19 HARV. J.L. & TECH. 1, 9 (2005).

² To be clear, we are not arguing for or against imposing a network neutrality rule on broadband network providers. Our view is that if Congress, the Federal Communications Commission, or some other authority imposes such a rule, it should allow network providers to take steps—including discriminating against certain kinds of traffic, applications, or protocols—to protect the trustworthiness of the network.

proposals for network neutrality rules, we show that advocates of network neutrality recognize the need to provide a trustworthiness exception to any neutrality obligation, but they differ in their prescriptions for the scope of this exception. We find three categories of exceptions: broad, medium, and narrow. Part II examines whether several plausible types of trustworthiness improvements would be permissible under these exceptions. We argue that the narrow trustworthiness exception prevents service providers from implementing trustworthiness improvements that are likely to be important in future networks; but the broad exception effectively swallows a neutrality rule. The medium exception avoids both of these problems. Still, getting the formal language of the exception right is only part what is necessary to establish a balance between neutrality and trustworthiness. Part III suggests that a trustworthiness exception provide some means to make ongoing assessments of whether network operators are using the exception appropriately. We propose the trustworthiness exception be conditioned on network providers' disclosure of trustworthiness-related discrimination.

I. Tracing Trustworthiness Through the Network Neutrality Debate

A. Defining Trustworthiness

A trustworthy system has been described as one that “does what people expect it to do—and not something else—despite environmental disruption, human user and operator errors, and attacks by hostile parties.”³ Trustworthiness is a “multidimensional” concept

³ National Research Council, Computer Science & Telecommunications Board, *Trust in Cyberspace* 13 (ed. Fred B. Schneider) (1999), <http://www.nap.edu/readingroom/books/trust/trustapk.htm> [hereinafter CSTB, *Trust in Cyberspace*]. See also Marjory S. Blumenthal, *The Politics and Policies of Enhancing Trustworthiness for Information Systems*, 4 COMM. L. & POL’Y 513 (1999).

encompassing “correctness, reliability, security . . . privacy, safety, and survivability.”⁴

Security, in turn, means resistance to attacks that “can compromise the secrecy, integrity, or availability of data and services.”⁵

Where the Internet is concerned, trustworthiness is important for a number of reasons. Computer networks have become elements of our nation’s infrastructures. Other highly developed nations are following suit. Network-based attacks, which can last for days, could have major effects on a national economy. For example, in May 2007, Estonia suffered a distributed denial of service attack that brought banking and other services to a halt for several days.⁶ Vulnerabilities in a network can also lead to leaks of personal information, potentially leading to a loss of privacy, personal financial losses, and revelations about candidates that might well alter the outcome of national elections.

To view network neutrality through the lens of network trustworthiness, concrete examples of trustworthiness properties will be helpful. By focusing on properties, and hence what must be guaranteed, we avoid limiting the discussion to the known, specific attacks of today. Attacks co-evolve with defenses, but trustworthiness properties one might expect from a network are independent of threats and the attacks they might employ.

To start, we refine a model typically used to describe relationships on the Internet. When considering trustworthiness, it is important to recognize that individual end users are not the only consumers of data services that networks carry; the subnets comprising

⁴ CSTB, *Trust in Cyberspace*, *supra* note 3, at 14.

⁵ *Id.*

⁶ John Schwartz, *Bit Wars: When Computers Attack*, N.Y. TIMES, June 24, 2007.

the Internet also exchange traffic with one another. These interconnections depend on peering and transport agreements, whose significance will become evident in Part II.

With a network's customer expanded to include subnets (as well as individual users and computers, we can list examples of network properties that are useful for building trustworthy networked information systems. For each property, we discuss the extent to which the current Internet architecture provides support.

Confidentiality. A sender might want a guarantee that data she sends are not intercepted or stored and then later accessed by unauthorized third parties. Such unauthorized access can be prevented by encrypting data, and the current Internet protocols allow this because they do not distinguish between encrypted and unencrypted data.⁷

Communications Privacy. In addition to preventing third parties from gaining access to the contents of a communication, a user might wish to prevent others from learning about the very existence of a communication. Guarding against disclosure of this kind of information involves limiting the dissemination of traffic logs and restricting access to packets in transit. Currently, network operators alone decide whether to keep logs of the traffic they carry; the Internet architecture does not provide users with a means to direct a network provider not to log traffic.

Integrity. One of the Internet's core networking protocols, the Transmission Control Protocol (TCP), implements a guarantee that data accepted by a receiver have not been

⁷ In practice, the strength of any guarantee against a confidentiality breach depends on other factors: the strength of the encryption algorithm, the sender's and recipient's key management practices, the trustworthiness of any certificate authority involved, and whether the encrypted data are dumped and decrypted offline. These factors are related to

corrupted while in transit. Each TCP header contains a field for a checksum, which is a (more or less) unique numerical coding of the bit strings comprising the header and data in a TCP packet.⁸ A receiver independently calculates the checksum of incoming data and compares it to the checksum that is carried in the packet. A difference in these two checksums indicates the data were corrupted in transit and causes the receiver to discard the packet. The sender would then retransmit that packet. Thus, packets not discarded are identical on the sending and receiving ends of a communication.

Availability. The current Internet architecture offers only limited guarantees concerning availability. Specifically, the Internet architecture provides guarantees that users who persist for long enough in attempting to communicate will be able to do so, aided (in part) by the multiplicity of routes that packets may take from sender to recipient. TCP enforces the availability guarantee by requiring the sender to repeatedly retransmit a packet until an acknowledgment packet has been returned to the receiving computer.⁹ However, this particular delivery guarantee does not imply that packet delivery is timely, and TCP delivers data on a best-effort, first-in-first-out basis. Network providers can therefore shape (i.e., delete or delay) traffic based on source, destination, and application type. Traffic-shaping decisions lie with network providers, so they are beyond the control of most users.

Outages, such as those caused by earthquakes or accidental severing of network cables, might cause traffic to take suboptimal routes and leave destinations unreachable;

cryptography rather than network design. The point is that the current Internet need not be changed to handle encrypted traffic.

⁸ W. RICHARD STEVENS, *UNIX NETWORK PROGRAMMING*, Vol. 1, 32 (2d ed., Prentice Hall PTR, 1998).

⁹ *Id.* at 32.

but the Internet’s current routing architecture keeps other hosts usable during such outages. Though network design might help to mitigate some environmental threats, it is unlikely to defend against all of them.

Though the current Internet’s availability guarantees are celebrated, the constant exposure of Internet-connected systems to attacks has led some to contemplate making future networks support a guarantee *against* receiving traffic from certain hosts, which we deem a guarantee of isolation.¹⁰ (The Internet does not provide such a guarantee.) We discuss in Part 2 how such a guarantee would relate to network neutrality.

Correctness. The Internet currently employs the domain name system (DNS) for translating between names that are easy to use and remember, such as www.whitehouse.gov, and the numerical IP addresses actually used for routing packets. The DNS is vulnerable to a variety of attacks that undermine network trustworthiness. For example, by compromising the DNS, attackers can redirect users’ packets to malicious sites. The malicious host might then impersonate the legitimate host, allowing attackers to collect usernames and passwords. This form of attack facilitates identity theft and the fraudulent use of personal information to commit financial crimes.¹¹ The Internet itself (or its successors) may provide facilities for higher-level queries, such as the search engine queries that have become many users’ primary means of navigating the Internet, as

¹⁰ David Clark calls the isolation guarantee that we describe in the main text a “negative availability” guarantee. See Clark, *infra* note 44, at 7. Clark uses this turn of phrase to portray a guarantee to not *receive* traffic as the opposite of the Internet’s “[positive] availability” guarantee to a party *sending* a packet—that that packet will be delivered. We find the term “isolation” more descriptive and use that term throughout our article.

¹¹ For more details about attacks on DNS, see Security Associates Institute, *Attacking the DNS Protocol – Security Paper* (Oct. 29, 2003), at http://www.rootsecure.net/content/downloads/pdf/sans_attacking_dns_protocol.pdf.

well as queries that allow programs to find services.¹² Compromises to these services could severely harm the trustworthiness of those networks.

The rest of Part I traces the role of trustworthiness in the development of network neutrality. We start with *Hush-a-Phone* and *Carterfone*, two cases that helped form the nondiscrimination norm that is central to the network neutrality principle.¹³ In particular, we demonstrate that trustworthiness served as a limiting principle on the nondiscrimination principle defined in this line of cases and regulations. That is, telephone network operators are generally forbidden from discriminating against content, devices, and applications *except* when discrimination is necessary to protect the trustworthiness of the network. We then show how policymakers and scholars imported this structure into proposed network neutrality rules, limiting the nondiscrimination obligation with an exception that applies when a network operator discriminates to protect trustworthiness.

B. Trustworthiness as a Limitation on Nondiscrimination in Common Carrier Regulations

A basic conception of trustworthiness emerged as a limiting principle in the judicial cornerstone of network competition policy, the D.C. Circuit’s 1956 decision in *Hush-a-*

¹² The DNS-based attack discussed in the main text is but one illustration of the security problems that arise from the difficulty of authenticating (i.e., “establishing the truth of some claim of identity”) the sender (or receiver) of a message on the Internet. National Research Council, *Toward a Safer and More Secure Cyberspace* 5-1 (June 2007). Consequently, developing stronger and more widely used authentication mechanisms is one approach under discussion to address a large number of Internet security ills. *See id.* at 3-5.

¹³ *See* Tim Wu, *Network Neutrality, Broadband Discrimination*, 2 J. ON TELECOMM. & HIGH TECH. L. 141, 142 (2005) (“The link between anti-discrimination regulations and network innovation are as old as the *Hush-a-Phone* and *Carterfone* decisions, which controlled AT&T’s efforts to destroy innovative network attachments.”).

Phone Corp. v. United States.¹⁴ Hush-a-Phone sold a telephone receiver attachment that reduced background noise present at the speaker's location and also prevented the speaker's voice from being heard by others in close proximity. AT&T and the Bell companies sought to ban use of the Hush-a-Phone device under a rule that forbade the "attachment to the telephone of any device not furnished by the telephone company."¹⁵ At the end of a lengthy proceeding to hear Hush-a-Phone's complaint against AT&T's application of this "foreign attachment" rule, the FCC found that the lower volume and distorted sound of a Hush-a-Phone user's voice effected a "public detriment" to the phone system and, on this ground, upheld the Hush-a-Phone ban.¹⁶ The *Hush-a-Phone* court, however, found that the FCC's own findings did not support its conclusion and ordered the Commission to reverse the ban of Hush-a-Phone devices.¹⁷ In doing so, the D.C. Circuit announced a broader principle, which forms part of the intellectual foundation of network competition policy: the device prohibition was an "unwarranted interference with the telephone subscriber's right reasonably to use his telephone in ways which are privately beneficial without being publicly detrimental."¹⁸ The court did not specify what a "public detriment" might be, but it clearly recognized the possibility that one user's attaching the wrong type of device to the phone network, or using a device in the wrong way, could degrade or disrupt phone service for others. That is, new devices must not threaten the trustworthiness of the phone system as a whole. The device at issue in *Hush-*

¹⁴ *Hush-a-Phone Corp. v. United States*, 238 F.2d 266 (D.C. Cir. 1956).

¹⁵ *Id.* at 267 (internal quotation omitted).

¹⁶ *Id.*

¹⁷ *Id.* at 269.

¹⁸ *Id.* at 269. *See also In re Use of the Carterfone Device*, 13 F.C.C.2d 420, 423-24 (1968) (referring to the statement in the main text as "the principle of *Hush-A-Phone*").

a-Phone did not pose such a threat. Nevertheless, preserving the trustworthiness of the phone network was integral to the *Hush-a-Phone* principle.

More than a decade later, the FCC considered whether the Carterfone device, which allowed a mobile radio user to connect to a party on the phone network, had a “material adverse effect upon use of the telephone system” when deciding whether to prohibit it.¹⁹ The FCC found that a device that provided “nonharmful interconnection” of a telephone system user to a user off the grid did not prevent AT&T from “carry[ing] out its system responsibilities,” including maintaining a reliable phone system. The *Carterfone* court prohibited AT&T from discriminating against a device—AT&T had approved its own interconnection device—unless the device caused harm to the telephone network. In other words, AT&T could not ban a potential competitor’s device while offering a device that posed the same risks to trustworthiness; if it wished to ban a device that threatened trustworthiness, it had to ban all similar devices.

The nondiscrimination rule announced in *Carterfone* was broad but not unlimited. The court explicitly stated that it was “not holding that the telephone companies may not prevent the use of devices which actually cause harm, or that they may not set up reasonable standards to be met by interconnection devices.”²⁰ The court also emphasized that AT&T “remain[ed] free to make improvements to the telephone system” and to revise standards for interconnection devices accordingly.²¹ Furthermore, in the wake of

¹⁹ *In re Use of the Carterfone Device*, 13 FCC 2d 420 (1968) [hereinafter *Carterfone*]. AT&T argued in the *Carterfone* proceeding that allowing the device to connect to AT&T’s network would “divide the responsibility for assuring that each part of the system is able to function effectively”—a duty that AT&T asserted it should be solely responsible for bearing.

²⁰ *Carterfone*, 13 F.C.C.2d at 424.

²¹ *Id.*

Carterfone, the FCC issued rules that established a testing and certification process for devices manufactured to connect to the telephone system, to ensure that they would not harm the network.²² Thus, *Carterfone* did not leave network providers powerless to ban devices that harm network trustworthiness of the network. Instead, *Carterfone* provided a limited trustworthiness exception to telephone network providers' general nondiscrimination obligations.

The FCC followed the *Hush-a-Phone* principle when computer connections to the phone network became common. In the *Second Computer Inquiry*, the FCC again affirmed *Hush-a-Phone*'s and *Carterfone*'s articulation of a "consumer right" to use the network "in ways [that] are privately beneficial without being publicly detrimental."²³

More recently, as the FCC and federal courts have removed broadband service providers from common carrier regulations that applied to the telephone system, the Commission has begun to revisit the relationship between network access and network trustworthiness.²⁴ In the midst of these regulatory shifts, former FCC Chairman Michael

²² These "Part 68" rules are codified in 47 C.F.R. Part 68. Of particular interest here is the definition of "harm" in these "Part 68" rules that is provided in these rules:

Electrical hazards to the personnel of providers of wireline telecommunications, damage to the equipment of providers of wireline telecommunications, malfunction of the billing equipment of providers of wireline telecommunications, and degradation of service to persons other than the user of the subject terminal equipment, his calling or called party.

47 C.F.R. § 68.3.

²³ See *In re Second Computer Inquiry*, 77 F.C.C. 2d 384, ¶ 142 (1980) (quoting and citing *Hush-a-Phone* and *Carterfone*) [hereinafter *Computer II*].

²⁴ Much of the complicated history of these developments is recounted in *National Cable & Telecommunications Association v. Brand X Internet Services*, which held that broadband service delivered via cable modem is an "information service," and hence not subject to the common carrier regulations that apply to a "telecommunications service." See 545 U.S. 967, 974-80 (2005) (describing the history of FCC regulations concerning access to communications as well as the particular proceeding that led to *Brand X*); *id.*

Powell articulated four “Internet Freedoms,” which include the freedom to use applications and attach devices of users’ choice.²⁵ Consistent with prior network access regulations, Chairman Powell bounded some of these freedoms with trustworthiness considerations. Specifically, “Freedom to Access” was subject to network providers’ “legitimate needs to manage their networks,” and the “Freedom to Use Applications” was subject to the qualification that they “will not disrupt the network.”²⁶

Though the FCC has used trustworthiness in a simple and consistent way, it has not articulated in detail how to distinguish a genuinely trust-related instance of discrimination from a spurious one. This elaboration might not have been necessary in the past; when the FCC complied with *Hush-a-Phone*, it might have been plausible to think of a single entity as owning a communications network and defending it against threats arising from the ends of that network. By the time Chairman Powell described the “Internet Freedoms,” however, the diversity of network ownership, the extent of network interconnections, the diversity of devices connected to networks might, and the ability of

985-1000 (explaining the Court’s decision to uphold the FCC’s classification of cable modem services). Shortly after *Brand X* was decided, the FCC classified broadband Internet service delivered via DSL as an information service. *See In re Appropriate Framework for Broadband Access to the Internet over Wireline Facilities*, 20 F.C.C. Rcd. 14853 (Sept. 23, 2005) [hereinafter FCC, *Wireline Order*]. For a brief history of all of these proceedings, see John Windhausen, Jr., *Good Fences Make Bad Broadband: Preserving an Open Internet through Net Neutrality* 8-12, Public Knowledge Working Paper, Feb. 6, 2006, at <http://www.publicknowledge.org/content/papers/pk-net-neutrality-whitep-20060206> [hereinafter Windhausen, *Good Fences Make Bad Broadband*].

²⁵ Michael K. Powell, *Preserving Internet Freedom: Guiding Principles for the Industry*, 3 J. TELECOMM. & HIGH TECH. L. 5, 11-12 (2004) [hereinafter Powell, *Preserving Internet Freedom*]. The four “Internet Freedoms” are: (1) Freedom to Access Content; (2) Freedom to Use Applications; (3) Freedom to Attach Personal Devices; and (4) Freedom to Obtain Service Plan Information.

²⁶ *Id.* at 11.

attacks to cross from one provider's network to another have made the notion of providers managing "their" networks somewhat simplistic.

C. Trustworthiness as a Limitation in Network Neutrality Rules

Legal scholars and policymakers have applied trustworthiness-by-exception, essentially without modification, to proposed network neutrality rules. Their proposals contemplate that network operators will discriminate against traffic exchanged with other providers' networks to protect trustworthiness, but their proposals differ significantly in scope.

Proposals from legal scholars avoided spelling out what trustworthiness threats warrant deviation from network neutrality, and avoided enumerating what mechanisms are permissible to defend against these threats. The most detailed scholarly proposal was offered by Professor Tim Wu, a leading proponent of network neutrality. He recognized the challenge that protecting network trustworthiness poses to neutrality: "Spam, viruses, junk mail and telemarketing are different names for problems that every information network faces. What this suggests is that network security must be taken seriously, but also cannot become a blanket answer to any scrutiny of carrier practices."²⁷ Wu advanced a model network neutrality statute that would make discrimination permissible to "limit[] . . . the distribution of computer viruses, worms, and . . . denial-of-service or other attacks."²⁸ This proposal, however, does not analyze specific network defenses, nor does

²⁷ Tim Wu, *Wireless Net Neutrality: Cellular Carterfone and Consumer Choice in Mobile Broadband* 27, New America Foundation Wireless Future Program Working Paper #17, Feb. 2007, at

http://www.newamerica.net/files/WorkingPaper17_WirelessNetNeutrality_Wu.pdf.

²⁸ Wu, *Network Neutrality*, *supra* note 13, at 170.

it offer any guidance for ensuring that network operators do not use the exception to discriminate for reasons other than trustworthiness.²⁹ Other scholars, less solicitous of trustworthiness, handled the neutrality-trustworthiness interface even more vaguely³⁰ or have offered fewer details about the contours of a trustworthiness exception.³¹

To get a better sense of how a trustworthiness exception might work, we turn to network neutrality legislation that has been introduced in Congress in recent years.³² The proposals can be grouped according to the breadth of the exceptions they allow: narrow,

²⁹ Wu offers only a broad principle: “[A]bsent evidence of harm to the local network or the interests of other users, broadband carriers should not discriminate in how they treat traffic on their broadband network on the basis of inter-network criteria.” *Id.* at 171.

³⁰ See, e.g., Brett M. Frischmann & Barbara van Schewick, *Network Neutrality and the Economics of an Information Superhighway: A Reply to Professor Yoo*, JURIMETRICS (2007); Brett M. Frischmann, *An Economic Theory of Infrastructure and Commons Management*, 89 MINN. L. REV. 917, 1011 (2005) (stating that network neutrality might prohibit placing “security and spam regulation measures” at the Internet’s core—even if efficient and effective—and that “this . . . may be one significant cost of sustaining an infrastructure commons”).

³¹ See Windhausen, *Good Fences Make Bad Broadband*, *supra* note 24, at 50-51 (stating that a network neutrality rule should allow a network operator to “block[] spam, viruses, or threats to national or network security”).

³² Congressional forbearance from imposing a nondiscrimination obligation would likely leave service providers with broad power to block or degrade communications for security purposes without regard to their source or contents. See, e.g., Yoo, *Beyond Network Neutrality*, *supra* note 1, at 9, 22, 31, 71. In the absence of a Congressional network neutrality mandate, the regulatory levers that would remain to address discrimination by service providers include conditions on telecommunications provider mergers and FCC rulemakings. For discussions of the possibility of FCC intervention outside of the merger context, see *Wireline Order* ¶ 102, *supra* note 24 (reserving possibility that the FCC will use Title I ancillary jurisdiction to regulate broadband Internet access) and Harold Feld, *DSL Item Released—Coulda Been Worse*, WETMACHINE, Aug. 5, 2005, <http://www.wetmachine.com/totsf/item/333>. In the merger context, the FCC imposed a condition of “maintain[ing] a neutral network and neutral routing” on the merger of AT&T and BellSouth, effective for 30 months after closing. See Press Release, *FCC Approves Merger of AT&T and BellSouth Corporation* 8, Dec. 29, 2006, at http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-269275A1.pdf. For a proposal to give the FCC “antitrust-like” authority to adjudicate complaints about service providers abusing their market power, see generally Robert D. Atkinson & Philip J. Weiser, *A “Third Way” on Network Neutrality*, May 30, 2006, at

broad, and medium. For the remainder of this article, we use the term “trustworthiness exceptions” to refer specifically to these three classes. By “breadth,” we refer to conditions under which an exception would allow a network operator to discriminate as well as what it would allow a provider to discriminate against, e.g., traffic sources or destinations, applications, or protocols. We define the trustworthiness exceptions classes as follows:

- **Narrow Exception:** A network provider may not discriminate against traffic, applications, or protocols to protect trustworthiness. As we read bills that provide such an exception, a network provider would be limited to offering trust-related services, such as spam filtering or virus protection, so long as individual users may opt out of them.
- **Medium Exception:** A network provider may discriminate against content, applications, or protocols to protect the trustworthiness of the network, but it may not take into account any affiliation (or lack thereof) with a content, application, or protocol provider when deciding whether to discriminate.
- **Broad Exception:** A network provider may discriminate against content, applications, or protocols, so long as it does so to protect the trustworthiness of the network.

The Narrow Exception comes from a bill in the current Congress introduced by Senators Byron Dorgan and Olympia Snowe.³³ Specifically, the Dorgan-Snowe bill

<http://www.itif.org/files/netneutrality.pdf>.

³³ Similar exceptions appear in state-level legislation in New York and Maine. The New York State Assembly is considering a network neutrality resolution, which provides this security exception:

creates an exception for “protecting the security of a user’s computer on the network of such broadband service provider, or managing such network in a manner that *does not distinguish based on the source or ownership* of content, application, or service.”³⁴ The assumption in the Narrow Exception is that network providers can only deploy defenses that protect “their” networks alone and may do so only provided it does not degrade connectivity based on the source of content or the application or service in use.

The Medium Exception derives from the Network Neutrality Act of 2006, introduced by Congressman Ed Markey.³⁵ The text of the Markey bill’s trustworthiness exception

Nothing in this section shall be construed to prevent a broadband or Internet network provider from taking reasonable and nondiscriminatory measures . . . to manage the functioning of its network to protect the security and to offer parental controls and other consumer protection measures of such network and broadband or internet network services if such management does not result in discrimination among the content, applications, or services on the network.

A. 3980-B § 243(2)(A), <http://assembly.state.ny.us/leg/?bn=A03980&sh=t>.

Similarly, a bill introduced in the Maine legislature would have mandated “nondiscriminatory access” but permitted a service provider to “[p]rotect the security of a user’s computer or provide services in a manner that does not distinguish the source of ownership of content, application or service.” *See* LD 1675, <http://www.mainelegislature.org/legis/bills/billtexts/LD167501.asp>.

³⁴ *See* Internet Freedom Preservation Act, S. 215, 110th Cong. § 2, <http://thomas.loc.gov/cgi-bin/query/z?c110:S.215.IS>: (emphasis added). [hereinafter “S. 215” or “the Dorgan-Snowe bill”]. The predecessor to this bill contained an identical exception. *See* S. 2917, 109th Cong., <http://thomas.loc.gov/cgi-bin/query/z?c109:S.2917.IS>. Along similar lines, the Internet Freedom and Nondiscrimination Act of 2006 would have allowed prioritization of certain types of data, so long as broadband service providers treated all providers of such data equally. This bill did not explicitly mention security. Instead, it contained a number of exceptions that might encompass network security. For example, § 3(c)(1) would have allowed a service provider “to manage the functioning of its network, on a systemwide basis, provided that any such management function does not result in discrimination”; and § 3(c)(4) explicitly allows a provider to “offer consumer protection services (such as parental controls), provided that a user may refuse or disable such services.” *See* H.R. 5417, 109th Cong., <http://thomas.loc.gov/cgi-bin/query/z?c109:H.R.5417.RH>.

³⁵ H.R. 5273, 109th Cong. [hereinafter “Markey bill”], <http://thomas.loc.gov/cgi-bin/query/z?c109:H.R.5273.IH>. Congressman Markey recently introduced a

support reading it more broadly—as stated in the definition of the Medium Exception above—than the Narrow Exception. Though the security exception in the Markey bill would require providers to use “reasonable and nondiscriminatory measures” to protect security,³⁶ the overall structure of the bill suggests that not all forms of discrimination are prohibited. Specifically, the line between permissible and impermissible discrimination appears to be whether a service provider takes into account the distinction between content or services that it (or an affiliate) provides, versus an unaffiliated provider: a network operator may “manage the functioning of its network, on a systemwide basis, provided that any such management function does not result in discrimination between content, applications, or services offered *by the provider and unaffiliated providers*.”³⁷ Finally, like the Narrow Exception, the Medium Exception allows a network operator to offer “consumer protection services” that might include trustworthiness guarantees, so long as subscribers may opt out of them.³⁸

Finally, the Broad Exception removes any nondiscrimination requirement, though the single-firm view of network security remained in place. The Internet Consumer Bill of Rights Act, introduced by Senator Ted Stevens, would have provided such an exception by allowing a network operator to “protect the security, privacy, or integrity of the network or facilities of such provider, the computer of any subscriber, or any service, including by (A) blocking worms or viruses; or (B) preventing denial of service

substantially revised bill, the Internet Freedom Preservation Act of 2008, H.R. 5353, 110th Cong., http://thomas.loc.gov/home/gpoxmlc110/h5353_ih.xml.

³⁶ Markey bill, *supra* note 35, § 4(b)(3).

³⁷ *Id.* § 4(b)(1) (emphasis added).

³⁸ *Id.* § 4(b)(4).

attacks.”³⁹ Note that the Broad Exception, as written in the Stevens bill, does not qualify a network operator’s right to discriminate with any consideration of affiliation between the provider and the target of discrimination.

Despite differences on the issue of discrimination for the purposes of improving network trustworthiness, trustworthiness exceptions (as well as the four Internet Freedoms⁴⁰) share a common approach to the increasing need for coordination among service providers: they ignore it. All of these proposals reflect a single-firm outlook on trustworthiness—service providers may decide when to act in the interests of securing the subnets they operate (or their subscribers’ computers), albeit with varying levels of immunity from the broader nondiscrimination requirements. Whether this silence precludes providers from coordinating on matters of trustworthiness, or what may be sensible guidelines for determining whether a provider’s actions are sufficiently protective of its subnet in the case of coordinated defenses, are questions we do not settle here.⁴¹ Still, this silence is worth noting, given the importance that coordinated defenses will play in improving network trustworthiness.⁴²

³⁹ This act was Title IX of the Communications Opportunity, Promotion, and Enhancement Act of 2006, H.R. 5252, 109th Cong. § 906, <http://thomas.loc.gov/cgi-bin/query/z?c109:H.R.5252.RS:>.

⁴⁰ See *supra* Part I.A.

⁴¹ Whatever these boundaries may be, the antitrust laws would provide *some* limit on the kinds of information that providers may share, as well as the purposes for which they may share it. Specifically, Sherman Act § 1, 15 U.S.C. § 1, prohibits agreements that unreasonably restrain trade; and sharing information about industry practices may sometimes run afoul of this law. See *Complaint, United States v. Professional Insurance Consultants Ins. Co.* (D.D.C. June 24, 2005) (Civil No. 1:05CV01272), *available at* <http://www.usdoj.gov/atr/cases/f209700/209728.htm> (alleging that actuarial consulting firms moved toward an industry standard of limitations on liability clauses by sharing competitively sensitive information about such clauses and their efforts to implement them independently).

II. The Implications of Trustworthiness Exceptions in Network Neutrality Rules

In this Part, we consider whether proposals to enhance network trustworthiness would be permissible under the three classes of trustworthiness exceptions introduced above in Part I. Two questions guide our analysis of the relationship between each exception and network neutrality:

1. What trustworthiness improvements are available without discriminating against traffic based on its source?
2. What is left of network neutrality's general nondiscrimination principle if network operators may discriminate against communications sources, applications, or services in order to enhance network trustworthiness?

Our discussion answers these questions in the context of three guarantees that would help to improve network trustworthiness. These three guarantees are by no means exhaustive, but they are sufficiently diverse to illustrate trust-enhancing mechanisms that would be permissible (or forbidden) under each class of trustworthiness exception. Part II.A examines a trustworthiness guarantee that might require service providers to agree not to exchange traffic. By contrast, the guarantee discussed in Part II.B would require providers to relinquish their right not to exchange traffic with each other. Finally, a provider acting unilaterally could implement the privacy guarantee given in Part II.C

Note that, in the past, concerns about potential § 1 liability for sharing security-related information have prompted Congress to propose an antitrust exception for sharing such information. *See* Cyber Security Information Act of 2000, H.R. 4246, 106th Cong., *at* <http://thomas.loc.gov/cgi-bin/query/z?c106:H.R.4246>. *See also* Center for Democracy & Technology, *Davis-Moran Cyber Security Information Act—H.R. 4246*, May 5, 2000, *at* <http://www.cdt.org/security/000504davismoran.shtml> (criticizing the antitrust exemption as “unnecessary”).

effectively. Examining this range of trustworthiness guarantees permits us to evaluate whether network neutrality trustworthiness exceptions accommodate the range of defenses available today and that appear to be promising for the near future.

A. Isolation from Unwanted Traffic

The Internet currently does not guarantee that a user will remain isolated from—i.e., will not receive traffic from—a specific set of hosts.⁴³ Yet, an isolation guarantee would be useful to defend against distributed denial of service attacks. Blocking traffic from certain hosts could also prevent the spread of viruses or worms from one host to another. Limiting the spread of this malware, in turn, could interrupt the formation of “botnets”—networks of compromised computers under the control of a remote attacker—which can then be used to launch distributed denial of service attacks, send spam, or store data that are useful in committing financial crimes.

Current defenses against this malware, are implemented predominantly at the edge of the network. Firewalls, for example, block traffic with specific characteristics; and anti-virus programs installed on individual PCs reduce the end-user’s risk of executing malicious software. These defenses, though helpful, have significant limitations. Authors of worms and viruses have become adept at crafting programs that evade detection. Furthermore, firewalls are usually ineffective against denial of service attacks, because the attacks saturate network resources near the edge or on the target host—even if a firewall prevents traffic from reaching the intended target, that host nevertheless remains unavailable if its link to the Internet is saturated by attack traffic.

⁴² See our discussion of this point in the Introduction.

The questions of whether networks should support and will support isolation guarantees are being debated by technologists and others who are considering future Internet designs.⁴⁴ Still, the basic contours are clear enough to discuss within the context of network access competition policy. Basically, guaranteeing isolation would require automatic detection of malicious traffic and the quarantine of infected hosts.⁴⁵ Detecting malicious traffic, in turn, might require service providers to exchange network data⁴⁶ as well as agreements *not* to exchange traffic. This, because certain kinds of attacks, such as distributed denial of service attacks, might be perpetrated using traffic whose packet-level characteristics are indistinguishable from legitimate traffic. Only when traffic observations from many points on the network are correlated could a picture of an attack emerge.⁴⁷

1. Isolation as a Consumer Service

Suppose an ISP offers to its subscribers a package of trustworthiness services relating to isolation, e.g., filtering traffic from botnets, worms, and viruses and blocking traffic believed to be part of a distributed denial of service attack. Only under the Narrow Exception—which would not allow blocking based on the source of network traffic—could this service be prohibited. As noted above, successful identification of certain kinds

⁴³ David D. Clark, *Requirements for a Future Internet: Security as a Case Study*, ver. 2.0, Dec. 3, 2005, at http://find.isi.edu/presentation_files/David_Clark-Security-Requirements-2.pdf.

⁴⁴ See *id.* 7.

⁴⁵ See *id.*

⁴⁶ See, e.g., Yinglian Xie et al. *Forensic Analysis for Epidemic Attacks in Federated Networks*, in *Proceedings of the 14th IEEE International Conference on Network Protocols (ICNP 06)* (2006) [hereinafter Xie et al., *Forensics in Federated Networks*], at <http://www.cylab.cmu.edu/files/cmucylab06014.pdf>.

⁴⁷ See *id.*; Mark Allman et al., *Fighting Coordinated Attackers with Cross-Organizational Information Sharing*, ACM SIGCOMM HotNets V (2006).

of attacks depends upon finding patterns in the source and timing of traffic; without the ability to discriminate on the basis of network traffic source, such mitigation would be ineffective. The Medium and Broad Exceptions, which permit at least some discrimination based on source for network security purposes, would allow network providers to engage in such blocking or degradation.

Still, the Narrow Exception permits service providers to offer “consumer protection services” so long as each user may refuse or disable the service.⁴⁸ This approach, however, would severely limit the effectiveness of blocking virus (or worm, or botnet) related traffic, because if any single user refuses to accept the service, then the network provider might have to handle all sorts of undesirable traffic, which itself does collateral damage to other sites. For example, if a user opts out of her provider’s denial of service protection service, then the provider would be obligated to deliver any denial of service attacks that have been launched against her site. But, when a denial of service attack occurs, the attack traffic makes the network unusable by other sites that share a network with this recalcitrant user, because the denial of service traffic destined to her site is also clogging the pipes these other sites use.

2. Isolation as Provider Policy

Two widely repeated observations about computer and network security might make an isolation service inadequate. The first observation is that end-users are reluctant to invest much in improving security. The second observation—as the example of denial of service attacks illustrates—is that the insecurity of one host on a network can harm end-users at another host. These observations are related: end-users do not fully internalize the

benefits of their investment in security and, conversely, any given user may be victimized by attacks launched from the “weakest link” in the network. A possible response from network providers is to block suspected worm, virus, and botnet traffic for *all* its subscribers. That is, instead of offering isolation guarantees as a separate service, the service provider imposes them by default.⁴⁹ The Medium and Broad Exceptions, however, would allow this approach.

⁴⁸ See Dorgan-Snowe bill, *supra* note 34, § 12(b)(3). We discuss the difficulties in this approach later in this section.

⁴⁹ It is difficult to predict which course network providers would take. On the one hand, ISPs are developing managed security services aimed primarily at large enterprise customers; thus, at least some service providers see managed security services as a potential new source of revenue. See Sarah D. Scalet, *Pipe Cleaners*, CSOONLINE.COM (July 1, 2007), at http://csoonline.com/read/070107/fea_pipecleaners.html (noting that “[f]or now, and maybe for the long run, companies like AT&T will have to continue to make careful decisions about what traffic they can safely delete without violating their service-level agreements with customers or overstepping their bounds as common carriers that just pass bits from left to right.”).

On the other hand, at least some network operators have taken aggressive, blanket action to block traffic. Once discovered by users, such measures are not popular. For example, in October 2007, Comcast began blocking (or, in Comcast’s terms, “delaying”) traffic associated with the peer-to-peer file sharing program BitTorrent. *F.C.C. to Look at Complaints Comcast Interferes With Net*, N.Y. Times, Jan. 9, 2008, <http://www.nytimes.com/2008/01/09/business/media/09fcc.html>. Comcast defended the practice as “reasonable network management.” Ryan Paul, *FCC to investigate Comcast BitTorrent blocking*, ARS TECHNICA (Jan. 8, 2008), at <http://arstechnica.com/news.ars/post/20080108-fcc-to-investigate-comcast-bittorrent-blocking.html>. Public outcry ensued after the blocking was discovered, the FCC opened an investigation, and consumer groups filed a complaint with the FCC. The relationship between Comcast’s justification—“network management”—and trustworthiness is unclear, and it is also unclear that Comcast blocked BitTorrent to defend against a threat to trustworthiness. Still, this example serves to illustrate that network operators face a potentially difficult choice between applying blanket policies, which hold promise in providing trustworthiness guarantees but could provoke user backlash, and making such blocking part of a strictly optional service.

An alternative approach is for multiple network providers to federate and exchange data about possible attacks.⁵⁰ The rationale for this federation is that smaller service providers administer smaller slices of the Internet’s address space; unlike backbone providers or large ISPs, these providers might not command a sufficiently wide view of the Internet to identify subtle threats.⁵¹ A last-mile ISP might also agree to share information with a backbone provider. The backbone provider, which handles a higher volume of traffic and is likely to have a more comprehensive view of Internet traffic than a last-mile ISP, would be able to provide the ISP with a broader view than the ISP could obtain on its own. Finally, two or more backbone providers might agree to exchange information about malicious traffic in order to provide their respective downstream customers—last-mile ISPs or large enterprise networks—with guarantees that malicious traffic will be suppressed.

Neither coordination among last-mile ISPs nor coordination between an ISP and one or more backbone providers is addressed in network neutrality security exceptions or in the network neutrality debate more generally. Network neutrality security exceptions are silent about the prospect of coordination among network access providers to implement isolation guarantees. As was the case with vertically integrated operations—whether performed as a service that a subscriber requests, or as a default policy of the service provider—the key from an implementation perspective is being able to block traffic based on source.

⁵⁰ See Xie et al., *Forensics in Federated Networks*, *supra* note 46, for a discussion of how this might work in practice.

⁵¹ See Scalet, *Pipe Cleaners*, *supra* note 49 (quoting Gartner vice president John Pescatore: “[I]t’s not just economies of scale . . . It’s that the carriers have access to information that the individual enterprise doesn’t.”).

3. Isolation Under the Broad Exception

A final consideration raised by the examples in this section is whether the Broad Exception, which would allow a network operator to discriminate arbitrarily in response to a trustworthiness threat, could swallow a network neutrality rule.

Attackers have methods to remotely install malicious software that evades both firewalls and anti-virus software. For example, users risk unwittingly downloading malicious software simply by viewing Web pages that have been corrupted by attackers.⁵² These threats pervade the Internet; accordingly, a service provider might be able to find justification for degrading the performance of an application or to degrade or block connections to specific hosts on the Internet. In the absence of any nondiscrimination obligation, the provider would be free to block sites or degrade applications of non-affiliated providers, even if its affiliates' offerings were equally risky. Thus, the Broad Exception might shelter provider conduct only incidentally related to Internet trustworthiness—but which may be motivated by reasons not related to any aspect of trustworthiness.

The Medium Exception avoids this. A network operator discriminating against certain content, applications, or protocols for (ostensibly) trust-related reasons while leaving the services of affiliated providers undisturbed might support an inference that the network provider has taken affiliations into account. As a result, the discrimination

⁵² See Niels Provos et al., *The Ghost In The Browser: Analysis of Web-based Malware*, in *Proceedings of the First Workshop on Hot Topics in Botnets (HotBots)* (2007) (demonstrating how malicious HTML and JavaScript can be used to cause a browser to download malicious software automatically to an end-user's computer—a so-called “drive-by download”).

would not be protected by the exception; instead, it would violate the network neutrality rule.

B. Availability and Integrity: Attribution of Path

Internet packet routing is currently beyond individual users' control. Once Internet communications leave a sender's last-mile ISP's network, they are carried by backbone providers until they arrive at the receiver's ISP.⁵³ These backbone providers exchange traffic under barter agreements in an unregulated market. As others have noted, peering agreements are responsible for various problems, including sub-optimal routing and a lack of investment in innovations to the Internet's core.⁵⁴ Indifference to the route between a sender and a receiver makes connections between end points resilient to failures of some subnets (by giving service providers license to update routes as needed), but it requires users to trust the routing infrastructure for the entire Internet. Two examples will illustrate that routing control by users would provide useful guarantees for improving network trustworthiness.

First, consider a user who trusts routers only in certain countries. For instance, this user might be a defense industry consultant who is traveling abroad and needs to

⁵³ See FTC, Staff Report, *Broadband Connectivity Competition Policy* 25-26, June 2007; Paul Laskowski & John Chuang, *Network Monitors and Contracting Systems: Competition and Innovation* 183, in *Proceedings of ACM SIGCOMM* (2006) [hereinafter Laskowski & Chuang, *Network Monitors*]; Syliva Ratnasamy, Scott Shenker & Steven McCanne, *Towards an Evolvable Internet Architecture* 315, in *Proceedings of ACM SIGCOMM* (2005); Ramesh Johari & John Tsitsilis, *Routing and Peering in a Competitive Internet*, Jan. 30, 2003 [hereinafter Johari & Tsitsilis, *Competitive Internet*]; David D. Clark, Karen R. Sollins, John Wroclawski & Robert Braden, *Tussle in Cyberspace: Defining Tomorrow's Internet*, in *Proceedings of ACM SIGCOMM* (2002)

⁵⁴ See Johari & Tsitsilis, *Competitive Internet*, *supra* note 53 (discussing "hot potato" routing under backbone provider peering agreements); Laskowski & Chuang, *Network Monitors*, *supra* note 53 (analyzing how peering agreements diminish incentives to invest in core Internet innovation).

communicate confidentially with her colleagues in the United States. But she surmises that her communications are likely to pass through countries that monitor the contents of Internet communications and would be highly motivated to try to break the encryption on communications relating to the U.S. defense industry.⁵⁵ If this user can control the routes that her communications take, she will be able to ensure that those communications travel only through countries whose routers she trusts; she would no longer have to trust the entire Internet.

A second example is a guarantee of *disjoint* paths, i.e., paths that do not rely on any of the same routers. The use of such paths increases the probability of delivering any given packet, because the probability of failure (or compromise) of a machine on any given path is independent of the other paths.⁵⁶

Providing stronger routing guarantees—whether a guarantee to follow a route specified by an end-user or a service provider’s guarantee of diverse routing—requires coordination among network access providers. Specifically, to implement these guarantees, network providers would have to: (1) implement a technical mechanism to express and communicate preferred routes; (2) agree to follow route specifications, and (3) provide some means for others to verify that a given provider had followed its promise to route traffic in the specific manner.⁵⁷

⁵⁵ Information about the likely route of an Internet communication can be obtained by using the `traceroute` command on Unix and Mac systems, or the `tracert` command on a Microsoft Windows system.

⁵⁶ An alternative to full user control over the routes for their communications is simply to provide guarantees of diverse routing. The current Internet architecture does not support these guarantees, either.

⁵⁷ For a proposal for how to implement these requirements in practice, see Karthik Lakshminarayanan et al., *Achieving Convergence-Free Routing Using Failure-Carrying*

We set aside the considerable change in economic relationships among last-mile and backbone providers that would be necessary to achieve such guarantees, in order to examine how they would be evaluated under a trustworthiness exception to network neutrality. In both examples we presented, end-users sought guarantees concerning paths their communications would take. The network provider did not draw distinctions among the end-points to which these users wanted to connect. In other words, a network provider's ability to offer attributions of path does not necessarily imply that the provider would use control over routing to degrade performance based on the end-user's choice of application or the identity of the other party to the communication. So long as the end-user controls this choice, these guarantees would fall within the scope of even the Narrow Exception.

A more difficult question arises if a network provider were to select routes based on its own security considerations. As a practical matter, a network architecture that provided routing guarantees could allow providers to discriminate against traffic based on its source or the application in use. A provider might decide, for example, that a particular Web browser leaves its users unacceptably vulnerable to the installation of malicious software by remote attackers. This vulnerability, the network provider might conclude, threatens the security of the network by opening it to further propagation of malicious software, or by enlisting it in distributed denial of service attacks.⁵⁸ Suppose that the

Packets, in *Proceedings of ACM SIGCOMM* (2007), at <http://www.cs.berkeley.edu/~mccaesar/papers/fcp.pdf>.

⁵⁸ We are aware that an ISP may have incentives to be disingenuous, using trustworthiness as a *pretext* to degrade service when the primary motivation may be a financial agreement with the provider of another, similar service. Indeed, the existence of such an agreement would create some suspicion about the service provider's motives. To keep this example simple, however, we assume that the service provider acts solely to

provider further reasons that alternatives to this browser with the same functionality are available at no cost. A network architecture that supports path attribution would allow the provider to choose relatively slow routes for requests from that browser, thus degrading the service based on the application that the subscriber has chosen.

In this case, the service provider would likely run afoul of the Narrow Exception. The service provider has clearly decided in this example to degrade the performance of a particular application, something that the narrow security exception flatly prohibits.⁵⁹

The Broad Exception, however, probably offers some cover for the service provider's decision to degrade the performance of the browser in question through route manipulation. The network provider in this example acted to preempt remote threats to the security of subscribers' computers by penalizing users who used a relatively vulnerable browser. This exception would allow a provider to block traffic from worms or viruses, or to "prevent[] denial of service attacks."⁶⁰ There is no requirement that the service provider act only to prevent or counter a denial of service attack once it is underway; a set of logically connected considerations—discouraging vulnerable browser use by degrading its performance might prevent malicious software installation, and thus prevent the use of such software to carry out denial of service attacks—might suffice to bring the provider's conduct within the scope of the Broad Exception. In addition, this

impose a penalty for using a highly vulnerable browser. We have much more to say about trustworthiness as a pretext for discrimination in Part III.

⁵⁹ See Dorgan-Snowe bill, *supra* note 34, § 12(1)(b) (requiring a service provider to manage security in a manner that "does not distinguish based on the source or ownership of content, *application*, or service") (emphasis added); Markey bill, *supra* note 35 (requiring a provider to protect of the security of its network or a subscriber's computer using "reasonable and *nondiscriminatory*" measures) (emphasis added).

⁶⁰ See H.R. 5252, *supra* note 39, § 906(1).

security exception would allow a provider to prevent “unauthorized” uses of its network, without any restrictions on the means employed to achieve that goal.⁶¹

To see the full implications of the Broad Exception, consider a slight change to our example. Suppose that alternative browsers have the same type of security vulnerability as the browser that the provider discriminates against. Since complex network applications such as Web browsers are almost certain to have at least some security vulnerabilities, network providers could cite network trustworthiness reasons to block or degrade traffic going to specific browsers, even if the primary motivation for such discrimination is to favor one or more browsers over others.

The Medium Exception would provide far less cover for using trustworthiness to limit a network neutrality rule. By forbidding network providers from considering their affiliation (or lack of affiliation) with application, content, or service providers when deciding whether to engage in trustworthiness-related discrimination, the Medium Exception would require a provider to articulate some trustworthiness rationale for discriminating against one browser while leaving others with similar security vulnerabilities unaffected.⁶²

In this regard, the Medium Exception essentially follows *Carterfone*’s trustworthiness rule, which allowed network providers to set standards for network trustworthiness and

⁶¹ *Id.* § 906(3). Note that other provisions of H.R. 5252’s security exception do not limit this exception. Users would have the right to run any application “without interference from an Internet service provider, *except as otherwise provided by law*” (emphasis added). *Id.* §§ 903(a)(7), (b)(1).

⁶² Determining whether the overall security of two applications is similar is not, in general an easy task; the disciplines of computer and network security are dogged by the difficulty of devising metrics to measure security. See National Research Council, *Toward a Safer and More Secure Cyberspace*, *supra* note 12, at 6-9-6-20.

ban devices that violate them, so long as the provider treats similar devices similarly.⁶³ If the network provider could produce such a rationale, the Medium Exception would permit the discrimination. Thus, the Medium Exception is flexible without allowing any claim of enhanced trustworthiness to protect discrimination from a network neutrality rule.

A potential objection to the Medium Exception is that it does not require the network operator to tailor its discrimination as narrowly as possible to address a given threat to trustworthiness. For instance, in the context of our browser example, the Medium Exception might allow a network provider to require the use of specific browser—a highly restrictive form of discrimination that could make it difficult for new applications to find users—when it would suffice to degrade or ban outright browsers that did not meet the provider’s security requirements.

Two considerations remove some of force of this objection. First, actual legislation, such as the Markey bill, which provided the impetus for our defining the Medium Exception, could limit network providers to *reasonably* discriminatory measures. It is beyond the scope of this article to suggest a standard for reasonableness, but it might include whether a network provider could have chosen less sweeping measures for achieving the same trustworthiness objective. Second, as we discuss in Part III, any trustworthiness exception to network neutrality should require network providers to justify the exception by disclosing relevant details about their use of discrimination. This disclosure requirement would allow users and regulators to monitor uses of discrimination, which might, in turn, lead providers to choose narrow defenses.

⁶³ See the discussion of *Carterfone* in Part I.B, *supra*.

C. Privacy and Confidentiality: Guarantees Against Logging

The conditions of network access—and the role of trustworthiness as a limitation on network neutrality obligations—encompass more than whether service providers degrade or block communications involving certain hosts or applications. To take one example, service providers play an essential role in setting guarantees of Communications Privacy.⁶⁴ In contrast to the trustworthiness guarantees discussed above, which individual service providers have relatively limited power to make, providers exert significant control over Communications Privacy guarantees. Competition among service providers shows promise to strengthen Communications Privacy guarantees, yet this dimension of competition is one that the network access policy debate has largely ignored.⁶⁵ This section expands the framework of network access competition to include end-user Communications Privacy.

Communications Privacy fits naturally into the framework that we have established for relating network trustworthiness to network access competition. First, Communications Privacy protection affects end-users' decisions about Internet use. For example, a user who is concerned about breaches of Communications Privacy might avoid visiting certain websites out of fear that her use will be revealed (or used in public

⁶⁴ As stated in Part II.A, Communications Privacy pertains to preventing third parties from learning of the *existence* of traffic to or from a party. This is but one element of privacy. It is nevertheless a focus appropriate for a discussion focused on the intersection between trustworthiness and network competition policy.

⁶⁵ There is some evidence to support the assertion that competition leads to increased individual information privacy. This has been most extensively explored in the context of search engines. A recent report from the Center for Democracy & Technology (CDT) finds, for example, that “search engines are now competing to provide the best privacy protections.” CDT, *Search Privacy Practices: A Work In Progress* 1, Aug. 2007, at <http://www.cdt.org/privacy/20070808searchprivacy.pdf>.

or private surveillance).⁶⁶ Thus, Communications Privacy guarantees could help to promote the goal of openness on the Internet that network neutrality advocates seek. As the Computer Science and Telecommunications Board of the National Research Council wrote in a recent report that makes “confidentiality of stored information and information exchange” part of a “Cybersecurity Bill of Rights”:

One central function of information technology is the communication and storage of information. Just as most people engage in telephone conversations and store paper files with some reasonable assurance that the content will remain private even without their taking explicit action, users should expect electronic systems to communicate and store information in accordance with clear confidentiality policies and with reasonable and comprehensible default behavior.

. . . As a particularly important way of ensuring confidentiality, responsible parties should have the technical capability to delete or expunge selected information that should not be permanently stored.⁶⁷

A second reason to view Communications Privacy as standing on equal footing with availability and integrity guarantees is that individual users—and technical approaches that focus on the edge of the network—are limited in what they can do to support Communications Privacy. Anonymizers provide some measure of Communications Privacy by making traffic analysis more difficult, but these technical measures can be

⁶⁶ See Helen Nissenbaum, *Privacy as Contextual Integrity*, 79 WASH. L. REV. 119, 121 (2004) (discussing how end-users’ online activities are recorded, stored, and analyzed into individual profiles for commercial use).

⁶⁷ National Research Council, *Toward a Safer and More Secure Cyberspace*, *supra* note 12, at 3-4.

cumbersome to use and do not address logging by ISPs.⁶⁸ Thus, like the isolation and path attribution guarantees discussed earlier, Communications Privacy guarantees could provide a basis for network service provider differentiation and competition in the near term and impetus for technical improvements in the longer term. In other words, technical and policy decisions about Communications Privacy will be made alongside decisions that affect other elements of trustworthiness as well as the Internet's support for innovation and openness.

The current Internet architecture does not provide technical guarantees to protect individual Communications Privacy. Last-mile ISPs, backbone providers, and Internet hosts (such as e-commerce sites) set their own network traffic logging policies. United States law does not require network access providers to retain data, but, on the other hand, it does not impose limits on the amount of data that these providers may retain.⁶⁹ Though details about data retention practices of specific network service providers are scarce, some prominent providers retain significant amounts of data about subscribers.⁷⁰

Thus, a straightforward and potentially far-reaching means of compromising individual Communications Privacy on the Internet is for a last-mile provider to link a

⁶⁸ See, e.g., Tor: Anonymity Online, Aug. 10, 2007, at <http://tor.eff.org/>. Tor uses a distributed network of servers to route communications in a manner that makes them resistant to traffic analysis by parties with access to network traffic logs.

⁶⁹ The Electronic Communications Privacy Act (ECPA) does establish a data *preservation* requirement; under specific circumstances, a service provider must preserve data that it has in its possession, but the ECPA has no provision to require retention prospectively. See 18 U.S.C. § 2703(f).

⁷⁰ See *Recording Indus. Ass'n of Am. v. Verizon Internet Servs.*, 240 F. Supp. 2d 24, 28 (D.D.C. 2003), *remanded by* 351 F.3d 1229 (D.C. Cir. 2003); *Charter Comms., Inc., Memorandum in Support of Motion to Quash Subpoena Served by Recording Indus. Ass'n of Am.*, (E.D. Mo., Oct. 3, 2003) (No. 4:03MC00273CEJ) (not arguing that Charter did not have the information necessary to comply with the RIAA's subpoena for personal identifying information linked to an IP address).

user's personal identifying information to his or her IP address and a list of addresses that subscriber visited. Whether a provider makes this link voluntarily or under compulsion,⁷¹ last-mile providers occupy a central role in setting communications privacy protections because they control subscriber information, IP address assignments, and may retain logs of subscribers' Internet use.⁷² Backbone providers may log Internet communications records but typically do not have the information necessary to link these records to individuals. Individual websites, on the other hand, may collect information about individuals but typically do not control the same breadth and volume of data that a last-mile a last-mile provider does. Thus, Communications Privacy guarantees from a last-mile provider, such as a policy limiting the scope and duration of data retention, could significantly reduce threats to privacy, though they would not eliminate them. This guarantee discussed here would gain little strength from coordination among different providers; so it lends itself to unilateral implementation by a single provider.

Raising the profile of Communications Privacy guarantees as a dimension of network provider competition would begin with seeking more information about current practices.⁷³ This, in turn, would provide end-users with sufficient information to

⁷¹ One of the three titles in the Electronic Communications Privacy Act (ECPA) regulates the circumstances under which a service provider may disclose such data voluntarily as well as in response to subpoenas or other compulsory process. *See* 18 U.S.C. §§ 2701-2710.

⁷² Though the details of specific network access providers' data retention practices are not publicly known, several sources of evidence suggest that they log considerable amounts of information about their customers' Internet use. And last-mile providers have a strong incentives—protecting against fraud, abuse, and bandwidth hogs—to keep information that will allow them to identify an IP address with a particular subscriber.

⁷³ Obtaining this information is likely to pose a significant challenge. *See* Ryan Singel, *Why ISP Data Survey Matters: One Smart Lawyer's Take*, THREAT LEVEL, Mar. 29, 2007, at http://blog.wired.com/27bstroke6/2007/03/why_isp_data_su.html (discussing the

discipline service providers in the marketplace, either by registering complaints with their providers or switching to a different one. Thus, improving Communications Privacy guarantees could follow naturally from greater disclosure of service plan information—which is a pillar of the current network neutrality regulatory environment—provided that policymakers, market participants, and advocates recognize that Communications Privacy as an element of trustworthiness that competition could help to improve.⁷⁴

D. Trustworthiness and Wireless Net Neutrality

A final test of trustworthiness as a limitation on network neutrality obligations concerns wireless networks. The rapid increase in the number of cell phones and other handheld devices that can access the Internet, as well as the FCC’s upcoming auction of “beachfront” spectrum, have focused attention on whether wireless networks should be subject to network neutrality principles.⁷⁵ A great deal remains unsettled in the FCC’s plans for wireless spectrum,⁷⁶ and we seek to remain agnostic about whether network neutrality is desirable for wired networks. Therefore, we do not take a position on wireless network neutrality, focusing instead on reasons to be cautious about extending network neutrality principles without modification into the wireless context.

difficulties involved in “ferret[ing] out how ISPs store and share user Internet usage histories”).

⁷⁴ See Powell, *Preserving Internet Freedom*, *supra* note 25 (discussing the “Freedom to Obtain Service Plan Information” as one of the four Internet Freedoms articulated by former FCC Chairman Michael Powell).

⁷⁵ The final rules governing the auction are in *In re Service Rules for the 698-746, 747-762 and 777-792 MHz Bands, Second Report and Order*, WT Docket No. 06-150, Aug. 10, 2007. For a brief overview of the auction, see Stephen Labaton, *Airwaves, Web Power At Auction*, N.Y. TIMES, Jan. 22, 2008, at C1.

⁷⁶ See Shaheen Samavati, *A New World of Wireless Is Just Around the Band*, CLEVELAND PLAIN-DEALER, Jan. 24, 2008, at C1 (describing possible uses of the spectrum being auctioned).

Wireless network operators, such as cellular phone carriers, have little ability to contain devices that threaten the trustworthiness of the network. Suppose, for example, that a cell phone has acquired software that causes it to send a flood of traffic, such as text messages, to other cell phones. Since one user's cell phone may interact with another before reaching provider-controlled equipment, cell phones are at greater risk for being made unavailable by their peers than are computers connected to the Internet by cable modem or DSL.⁷⁷ In contrast to cell phones, cable and DSL modems route all traffic through provider-controlled equipment, affording the provider an opportunity to contain individual devices that harm the availability of the provider's network.

A key question is whether differences between wired and wireless architectures warrant enacting a relatively narrow network neutrality rule for wireless networks. We would argue that differences in these architectures permit different trustworthiness threats, and that should be taken into account when settling on a network neutrality rule. It may turn out, after a full consideration of threats facing cell phones and other wireless devices, that the class of attacks involving one cell phone harming availability (or other trustworthiness properties) of another cell phone is unimportant compared to, say, Internet-based attacks on cell phones.⁷⁸ Still, policymakers and other participants in the network neutrality debate should not assume that all network architectures present identical considerations at the intersection of neutrality and trustworthiness.

⁷⁷ Recent work has shown that a relatively small amount of text message traffic can render cell phones unavailable for their intended uses. See William Enck et al., *Exploiting Open Functionality in SMS-Capable Cellular Networks*, in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)* (Nov. 2005), <http://www.smsanalysis.org/smsanalysis.pdf>.

⁷⁸ Cell phones might be particularly susceptible to such attacks. See *id.*

III. Keeping Trustworthiness Exceptions Limited Through Disclosure

Irrespective of the scope of trustworthiness exception enacted as part of a network neutrality rule, the question of how to enforce network provider compliance with the exception remains. Of particular concern is prevention of *pretextual* uses of the trustworthiness exception. In this Part, we address the potential for a network provider to assert that it is discriminating against traffic to protect network trustworthiness when, in fact, the discrimination does little or nothing to achieve that goal.

We consider in Part III.A three mechanisms to keep uses of a trustworthiness exception in check: leaving trustworthiness mechanisms unspecified and leaving enforcement entirely to ex post enforcement actions by users or regulators; writing the trustworthiness exception to specify all mechanisms permitted under the exception; and requiring network providers to disclose details regarding their uses of the trustworthiness exception. Of these, we conclude that the disclosure approach seems best suited for this job. We then set forth some additional considerations for trustworthiness-related disclosures in Part III.B.

A. Mechanisms to Deter Trustworthiness-as-Pretext

The simplest approach to enforcing compliance with a trustworthiness exception is to allow government enforcement agencies or end users to sue network providers over bogus uses of the exception, but to refrain from adding anything to the network neutrality rule that would aid enforcement. Discrimination that does not fall within a trustworthiness exception (or other statutory exception to the nondiscrimination obligation) is deemed

impermissible discrimination, and the conduct is subject to whatever penalties the network neutrality rule provides.

Experience with network operator practices that have animated the network neutrality debate give some support to this approach; several instances of outright blocking have been discovered, brought to public attention, and, for the most part, quickly reversed by the network operator.⁷⁹ Indeed, this is the approach taken by most of the legislation reviewed in Part I.

There are two potential problems with this non-regulatory approach. First, though some network neutrality violations will be readily apparent to users, others might be more subtle. Relying on user vigilance does not suffice. Second, leaving the trustworthiness exception without any structure to support its application leaves network providers with a stark choice: implement trustworthiness-enhancing discrimination in a way that a court finds to fit the exception, or face liability for violating the neutrality rule. This choice might cause network operators to err on the side of avoiding liability at the cost of trustworthiness. Conversely, users would also face a binary choice if they discover that their network operator is discriminating in an unacceptable fashion: sue (if private suits are allowed) or change providers. To rely on the latter would assume the market for network access is competitive enough to allow user behavior to discipline network operators' conduct. In other words, it would assume away the problem that has motivated the network neutrality debate in the first place.

A second approach to enforcing compliance would be to enumerate all permissible instances of pro-trustworthiness discrimination. To provide sufficient guidance, this

enumeration would have to go beyond listing the *threats* against which discrimination is permissible.⁸⁰ It would have to leave the means of addressing trustworthiness threats completely unspecified, allowing network operators to choose draconian mechanisms that have effects that go well beyond defending against specific threats. Specifying not only which threats warrant departures from neutrality but also how a network operator may defend against these threats would address this problem. The approach, however, is likely to be extremely costly—if not impossible—to develop and inflexible in practice. As many of the examples in Part II illustrate, trustworthiness threats and guarantees are evolving; even if it were possible to specify all today’s threats and defenses, the taxonomy would quickly become outdated. And unless the trustworthiness exception kept pace with defenses, a listing of permissible defenses would fall behind.

Third, policymakers could add a disclosure requirement to the trustworthiness exception, requiring network operators to report each instance of a neutrality departure taken to protect network trustworthiness. This approach avoids the rigidity of specifying ex ante all circumstances in which discrimination is permissible. Instead, network operators decide for themselves when a particular threat warrants discrimination. Disclosure would allow either regulators or users to monitor use of a trustworthiness exception, thus reducing the risk of abusing the exception to create a market advantage. By having access to data from a wide variety of providers, enforcers would be better able to judge which practices are common in the industry, and thus would be less likely to

⁷⁹ See, e.g., Adam Liptak, *Verizon Blocks Messages of Abortion Rights Group*, N.Y. TIMES, Sept. 27, 2007, at A1.

⁸⁰ This was the approach taken in the Stevens bill, which would allow network operators to “block[] worms and viruses” and “prevent[] denial of service attacks.” H.R. 5252, *supra* note 39, § 906(1).

reflect one provider's favoring a content source or application provider. This, in turn, might lead to more nuanced assessments of provider conduct, reducing the risk that the threat of litigation would chill deployment of new network defenses.

To be effective, a disclosure requirement would need to be accompanied by a penalty for a provider's failure to comply. Otherwise, providers might determine that negative customer reaction or other risks (discussed below) outweigh their obligation to report. Auditing network operators' use of the exception by comparing an operator's internally documented and publicly disclosed instances of discrimination could provide a way to ensure compliance. These audits, however, would be invasive and potentially far more costly than the incidence of misreporting warrants. A mechanism with lower costs would be to condition a network provider's use of the trustworthiness exception as a defense to an alleged neutrality violation upon disclosing the discrimination at issue. That is, a provider could invoke the trustworthiness exception to defend discrimination only if it had disclosed the instance(s) of discrimination cited in an action brought under the network neutrality rule.

B. Striking the Right Balance for Trustworthiness Disclosures

Mandatory disclosures under the trustworthiness exception must balance several considerations. First, there is a trade-off between providing sufficient data to assess the merits of trustworthiness-related discrimination, on the one hand, and, on the other, not providing data that would put a network provider at a competitive disadvantage.

Second, disclosures could put the network operator at risk by revealing details about its network configuration and defenses.

Third, there is a risk that disclosing information about discrimination would reveal previously unknown vulnerabilities, leaving other providers' networks at greater risk of attack.

Fourth, a disclosure could reveal information relevant to areas other than trustworthiness, or to network neutrality for that matter. For example, suppose network providers *A* and *B* exchange traffic under an interconnection agreement that requires both providers to handle each other's traffic at least as favorably as they handle traffic from other providers. Suppose further that *A* discloses that it blocked botnet command traffic originating in *B*'s network. If *B* disputes *A*'s conclusion, it might use the information in *A*'s disclosure—which *B* might not otherwise have obtained—to sue *A*. In other words, the disclosure requirement might create a thicket of competitiveness, security, and contractual headaches for providers that are merely attempting to preserve their right to invoke the trustworthiness defense.

Limiting disclosure to the agencies that have the authority to enforce the network neutrality rule—most likely the FCC and the FTC⁸¹—would mitigate the problems discussed above. Limited disclosure would allow, for example, vendors of vulnerable products time to develop patches. It would also leave parties wishing to bring action against a network provider, for a network neutrality-related claim or otherwise, in the same position as they would be in without the disclosure requirement. Keeping these reports confidential would require either a specific exemption from the Freedom of Information Act (FOIA) or successful application of FOIA's exemption of trade secrets

⁸¹ In a recent report, the FTC made the case for its competence and authority under FTC Act § 5 to prosecute violations of network neutrality. *See* Federal Trade Commission, Staff Report, *Broadband Connectivity Competition Policy* 129-137 (June 2007).

and confidential commercial information.⁸² An agency could bolster the case for a FOIA exemption by promising the disclosing provider an instance of discrimination confidentiality.⁸³ This approach is in tension with the overall thrust of network neutrality, which seeks to provide greater transparency into network providers' practices. Agencies receiving discrimination reports could reduce this tension by providing public reports that de-identify the source of the information and remove information about specific threats and defenses.⁸⁴

A related question, whose answer may shed some light on appropriate recipient(s) for trustworthiness-related disclosures, is what information constitutes an adequate disclosure. A wide range of attacks and defenses might be the subject of disclosures,⁸⁵ making general guidelines difficult to develop. Although that task is too technical for the present article, we do point out that the agencies likely to enforce a network neutrality—the FCC and the FTC—have the capacity to develop detailed guidance, or even rules, to

⁸² See 5 U.S.C. § 552(b)(4).

⁸³ There is precedent for the FCC litigating FOIA requests for data of potentially similar sensitivity. See, e.g., *Center for Public Integrity v. FCC*, Memorandum Opinion, Case 1:06-cv-01644-ESH (Aug. 27, 2007) (holding that detailed data from network providers regarding broadband connections in particular geographic areas are exempt from disclosure by the FCC under 5 U.S.C. § 552(b)(4)).

⁸⁴ Determining exactly what information to put into a public report might be subtle. The problem of de-identifying data, for example, does not have any general solutions; and several high-profile cases had data re-identified after release. See Bruce Schneier, *Anonymity and the Netflix Dataset*, SCHNEIER ON SECURITY, Dec. 18, 2007, http://www.schneier.com/blog/archives/2007/12/anonymity_and_t_2.html (commenting on a paper that reported an algorithm that could uniquely identify 99% of anonymized Netflix movie reviews from eight such reviews and other publicly available data); Saul Hansell, *AOL Removes Search Data on Vast Group of Web Users*, N.Y. TIMES, Aug. 8, 2006, at C4 (discussing fallout from AOL's release of (weakly) de-identified search query histories of several hundred thousand users, which reporters and others promptly re-identified) Cite AOL search query stories; Netflix "anonymous" dataset re-identification.

⁸⁵ As discussed in *supra* Part II.

specify disclosure standards.⁸⁶ Both the FCC and FTC regularly conduct rulemakings on highly complex, technical topics. They also employ scientists capable of helping evaluate technical claims that would arise in this context. Moreover, hearings or rulemakings would allow public participation to help determine the kinds of data that network providers should submit to enforcement agencies. Though these reports might not be publicly available, this process would at least allow members of the public to ensure that regulators are receiving data helpful in assessing network providers' deviations from network neutrality to protect network trustworthiness.

Conclusion

This paper offers a few conclusions that, we hope, can advance the debate regarding network neutrality and network trustworthiness. Cyber threats are an increasingly urgent matter for network operators and end-users. A trustworthiness exception that does not allow a provider to discriminate based on the source of Internet communications is unlikely to give service providers sufficient latitude to respond to modern-day threats. And even if a trustworthiness exception allows service providers to offer security services, based on discriminating against traffic sources or application separately from basic Internet service, such a provision might leave providers incapable of protecting users from the “public detriments” that have set limits on the extent of network openness ever since *Hush-a-Phone* was decided.

⁸⁶ See Deirdre K. Mulligan & Aaron K. Perzanowski, *The Magnificence of the Disaster: Reconstructing the Sony BMG Rootkit Incident*, 22 BERKELEY TECH. L.J. 1157, 1224-30 (2007) (arguing that the FTC is a “natural place” to undertake the technically complex work of setting computer security-related policy); Pamela Samuelson & Jason Schultz, *Should Copyright Owners Have to Give Notice of Their Use of Technical Protection Measures?*, 6 J. ON TELECOMM. & HIGH TECH. L. 41, 69-70 (2007).

Still, a trustworthiness exception that does not impose any limits on discrimination could swallow the neutrality rule. The threats that currently face the Internet are far more varied and complex than those facing the telephone system in *Hush-a-Phone* and *Carterfone*, but we argued in Part II that the broad exception would allow at least some spurious claims of protecting security to serve as cover for practices that have, at most, a tenuous connection to network trustworthiness. To the extent that policymakers are concerned about service providers using a trustworthiness exception to evade a neutrality obligation, they should consider the reporting requirements discussed in Part III.

We have also identified the possibility that information sharing among providers to improve network trustworthiness could threaten competition that the network neutrality debate has thus far ignored. Determining whether these agreements could affect competition among network providers is an important area for future work, and it will require combining the findings of technical research with a more detailed empirical picture of the economic relationships among network providers and economic and legal theories for evaluating competition under these conditions.